


CONNEXIONS



The Interoperability Report

February 1992

Volume 6, No. 2

*ConneXions —
The Interoperability Report
tracks current and emerging
standards and technologies
within the computer and
communications industry.*

In this issue:

Archie.....	2
ISDN as an Open Platform...	10
Errata.....	15
NREN Bill signed into Law..	16
Book Reviews.....	19
Reader Feedback.....	20
Announcements.....	21

From the Editor

The notion that Internet applications comprise only the “major three” (Telnet, FTP, and E-mail) is rapidly being dispelled. In addition to the basic applications, a user needs *tools* to tell him *what* service/information is available, *where* it is available, and *how* to go about accessing that service. In past issues we’ve described Internet Library Systems, Wide Area Information Servers (WAIS), Resource Discovery, and other mechanisms. This month we look at yet another tool, namely *archie*. *Archie* is a resource indexing and discovery system that allows users to locate and identify files stored at anonymous FTP archive sites throughout the Internet. The article is written by two of *archie*’s principal architects, Peter Deutsch and Alan Emtage from McGill University.

The *Electronic Frontier Foundation* (EFF) is a public interest organization established in 1990 to educate the public about the democratic potential of new computer and communications technologies. EFF works to develop and seeks to implement public policies to maximize freedom, competitiveness, and civil liberty in the electronic social environments being created by these new technologies. In October 1991, the EFF testified before the House Energy and Commerce Committee, Subcommittee on Telecommunications and Finance in hearings regarding Telecommunications Infrastructure Legislation and Proposals. A summary of this testimony can be found starting on page 10.

Last summer we reported on the progress of legislation to provide for the *National Research and Education Network* (NREN). In December 1991, the bill was finally signed into law by President Bush. We asked Mike Roberts to once again give us an update from the nation’s capitol. His article appears on page 16.

Also in this issue are a couple of book reviews, several announcements, and one reader’s reaction to the article “Is Resource Discovery Hacking?” As always we welcome your comments, suggestions and questions, see page 23 for information on how to contact us.

Plans for INTEROP 92 Spring (May 18–22, Washington, DC) are already well underway. By now you should have received the Advance Program, if not call 1-800-INTEROP or 415-941-3399 to request your copy. There are still openings for *Birds of a Feather* (BOF) sessions. BOFs give you and your colleagues an opportunity to discuss network computing topics in an informal, after-hours atmosphere. If you’d like more information on BOFs or would like to organize one, please contact me at 415-962-2515 or send e-mail to ole@interop.com.

ConneXions is published monthly by Interop, Inc., 480 San Antonio Road, Suite 100, Mountain View, CA 94040, USA. 415-941-3399. Fax: 415-949-1779. Toll-free: 1-800-INTEROP.

Copyright © 1992 by Interop, Inc.
Quotation with attribution encouraged.

ConneXions—The Interoperability Report and the *ConneXions* logo are registered trademarks of Interop, Inc.

ISSN 0894-5926

The *archie* System: An Internet Electronic Directory Service

by Peter Deutsch and Alan Emtage, McGill University

Introduction

The existence of global, Internet accessible information providers has spurred the development of a number of mechanisms for locating and exchanging information. Distributed file systems, on-line file archiving mechanisms, electronic mail and bulletin boards, and even expert systems for locating and accessing technical expertise, are all services that are now available.

The huge size (and continued rapid growth) of the Internet offers a particular challenge to systems designers and service providers in this new environment. Before a user can effectively exploit any of the services offered by the Internet community or access any information provided by such services, that user must be aware of both the existence of the service and the host or hosts on which it is available. Adequately addressing this "resource discovery problem" is a central challenge for both service providers and users wishing to capitalize on the possibilities of the Internet.

In this article we describe *archie*, an Internet resource indexing and discovery tool that currently allows users to locate and identify files stored at anonymous FTP archive sites throughout the Internet. We provide a brief introduction to the system, along with an overview of the system architecture.

What is the *archie* service?

The *archie* service is a collection of resource discovery tools that together provide an electronic directory service for locating information in an Internet environment. Originally created to track the contents of anonymous FTP archive sites, the *archie* service is now being expanded to include a variety of other on-line directories and resource listings.

Currently, *archie* offers access to two databases. The first, the Files Database, contains the names of files located at anonymous FTP archive sites. The second, the *Whatis* database, contains the names and descriptions of thousands of software packages and other information on the Internet.

Users can search these databases using a variety of search and access methods by contacting an *archie* server. Access can be made using interactive sessions (provided they have a direct Internet connection) or through queries sent via electronic mail messages (provided they can at least gateway electronic mail messages onto the Internet).

Interactive access to *archie* may be through a conventional *telnet* session to a machine running an *archie* server or through a program that has been integrated into a larger system, such as the *Prospero* network distributed file system. Additional stand-alone clients are now being tested and are available over the network.

Why use *archie*?

The existence of the *archie* service allows those seeking information maintained by an *archie* server to limit their network search to a set of questions to a known server. The responses in turn offer pointers to specific Internet service providers. Once the existence and location of specific information or services has been determined using *archie*, traditional networking tools can be used for final access.

Programs have already been created that integrate an *archie* client with the *ftp* file transfer program or into larger information access services. This allows a user to first locate and then access information from *archie* sites using a single program.

The *archie* service today

Currently, *archie* tracks the contents of nearly 900 anonymous FTP archive sites containing some 1,600,000 files throughout the Internet. Collectively, these files represent well over 90 Gigabytes (90,000,000,000 bytes) of information, with additional information being incorporated daily. Anonymous FTP archive sites offer software, data and other information that can be copied and used without charge by anyone with a connection to the Internet.

The *archie* server automatically updates the listing information from each site about once a month, ensuring users that the information they receive is reasonably timely, without imposing an undue load on the archive sites or network bandwidth.

The *Whatis* database now contains some 3,500 package descriptions. These entries are not yet automatically maintained in the way the Files database is, but this is planned for a future release, as are additional databases for other collections of frequently changing information.

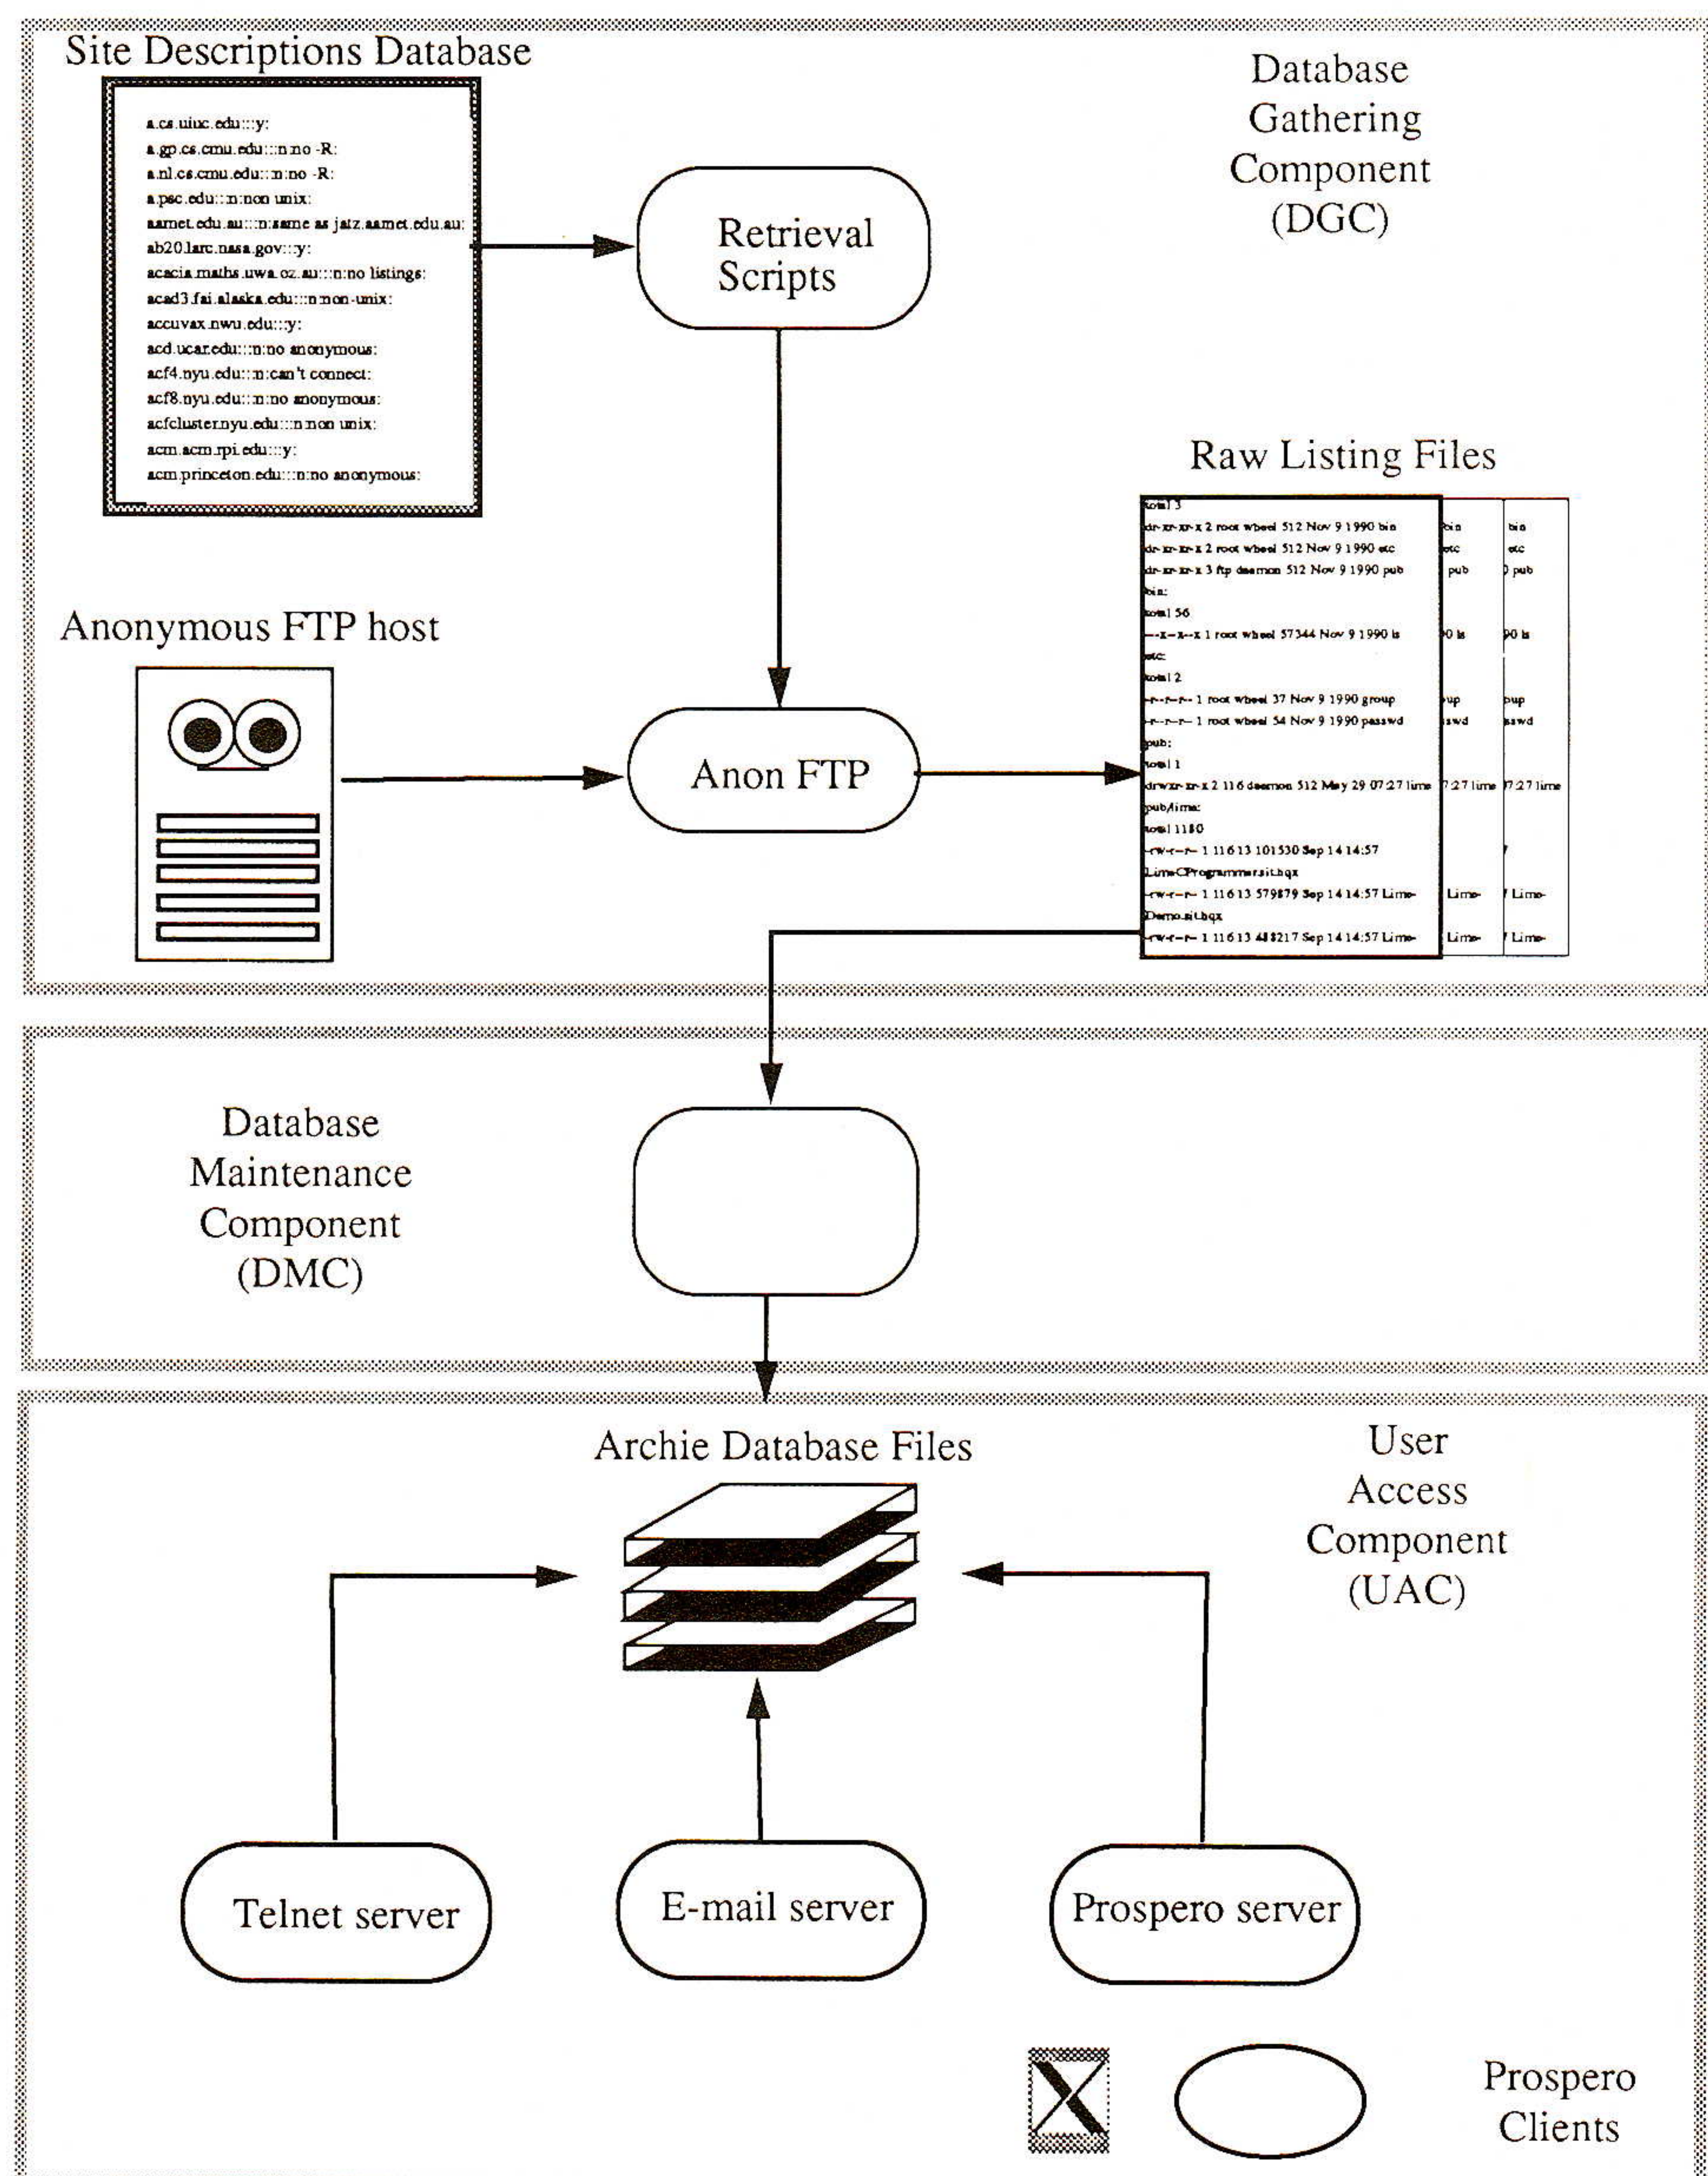


Figure 1: *Archie* system overview

continued on next page

The *archie* System (continued)

There are now nine *archie* servers in operation on the Internet. These include five general access servers (*archie.mcgill.ca*, *archie.funet.fi*, *archie.au*, *archie.sura.net*, and *archie.ans.net*), plus four more that are reserved for local users to reduce network bandwidth. These are located in Israel, Japan, New Zealand and the United Kingdom. Additional servers are planned for the coming months.

Architectural overview

The *archie* system offers the user a simple model for building, maintaining, and accessing a set of information databases in an Internet environment (see Figure 1). The entire system currently consists of only three major sub-systems, including the *Data Gathering Component* (DGC), the *Database Maintenance Component* (DMC), and the *User Access Component* (UAC).

Currently, the DGC and DMC are used only to maintain the Files Database. The *Whatis* database is maintained entirely by hand, although both are accessed through the same UAC channels. This is regarded as a serious shortcoming and will be changed in a coming release.

Data Gathering Component

The *Data Gathering Component* (DGC) is a fairly simple subsystem, consisting of standard UNIX shell scripts that are executed every 24 hours using the UNIX *cron*(1) facility. These scripts are used to connect to each monitored site, in turn, to fetch a recursive listing of the site's contents. This information is written to a "raw site listing" file on the *archie* server host, one for each site tracked.

Any number of strategies could be used to control the frequency of site updates. On the prototype *archie* server in Montreal we use a simple round-robin algorithm, cycling through the entire list of sites about once a month. This scheme was chosen for its simplicity and to assure site administrators and network powers-that-be that the *archie* system would not constitute an unwarranted drain on their resources.

Unfortunately, this simple updating strategy also means that some site information in the Files Database will be as much as 30 days out of date. Fortunately, few sites actually undergo radical change from month to month. Since a 30 day update cycle corresponds to 15 day average latency for the database as a whole this has proved acceptable in practice.

Other *archie* servers use more complicated updating schemes. As an example, the Australian server operated by AARNet currently cycles through all Australian archive sites every night, but tracks overseas sites by mirroring them synchronously from Montreal. This reduces the load somewhat on the heavily used transpacific link yet assures timely tracking of changes to those sites most visited by their users. The Australians also mirror a number of the most popular archived files onto a local archive to reduce the need for transoceanic access.

Since the *archie* site updating algorithm is implemented using a simple shell script and the *cron*(1) utility, changing the scheduling algorithm or frequency of updates is a relatively straightforward procedure, and we have seen a variety of techniques in this area among the *archie* site administrators. Balancing the need for accuracy and currency in the databases, the cost of gathering data and the cost of performing database site updates is currently done using empirical estimates and our previous experience with the system. We believe it is an area that would benefit from further study.

Site Description Database

Now that there are multiple *archie* sites, we will also have to begin addressing the issue of maintaining consistency between *archie* servers, with the twin goals of ensuring accuracy of the multiple databases and minimizing network bandwidth. We are currently working with individual *archie* site administrators to investigate mirroring and update strategies. Work in this area is expected to continue.

Operation of the DGC is controlled through a *Site Descriptions Database* (SDD) that lists each anonymous FTP site that we have discovered, along with additional information such as operating system in use at that site, whether a site is capable of providing a usable *ls -lR* file, commands to issue to the *ftp(1)* session during the fetch and whether we are currently tracking this site.

Currently the SDD is maintained entirely by hand. This is one of a number of pieces of *archie* that would benefit considerably from automation as time and resources permit. We also plan to make such site information available as an additional *archie* database in a future release. This would include a description of the site, access and storage policies, etc. Users would be able to search these entries in a manner similar to the *Whatis* database entries.

Discovering new sites

As there is still no generalized resource discovery or registration mechanism on the Internet, we continue to rely on site administrators or users to report new sites to us. This has become easier as we have become better known and we currently are aware of some 1,200 anonymous FTP sites on the Internet, although due to difficulties in obtaining usable site listings we currently actively track only about 900 of these sites.

The DGC uses proactive data gathering to ensure the internal accuracy and consistency of the *archie* Files Database. By automating the data gathering step, we provide the database maintenance component with information that has been verified accurate and in a suitable format. By periodically repeating this data gathering step we ensure database accuracy over time.

Database Maintenance Component

The *Database Maintenance Component* (DMC) is responsible for verifying the consistency of the raw site listings files and converting them into a format suitable for entry into the Files Database. This is currently done by three programs, each operating on one site at a time (see Figure 2).

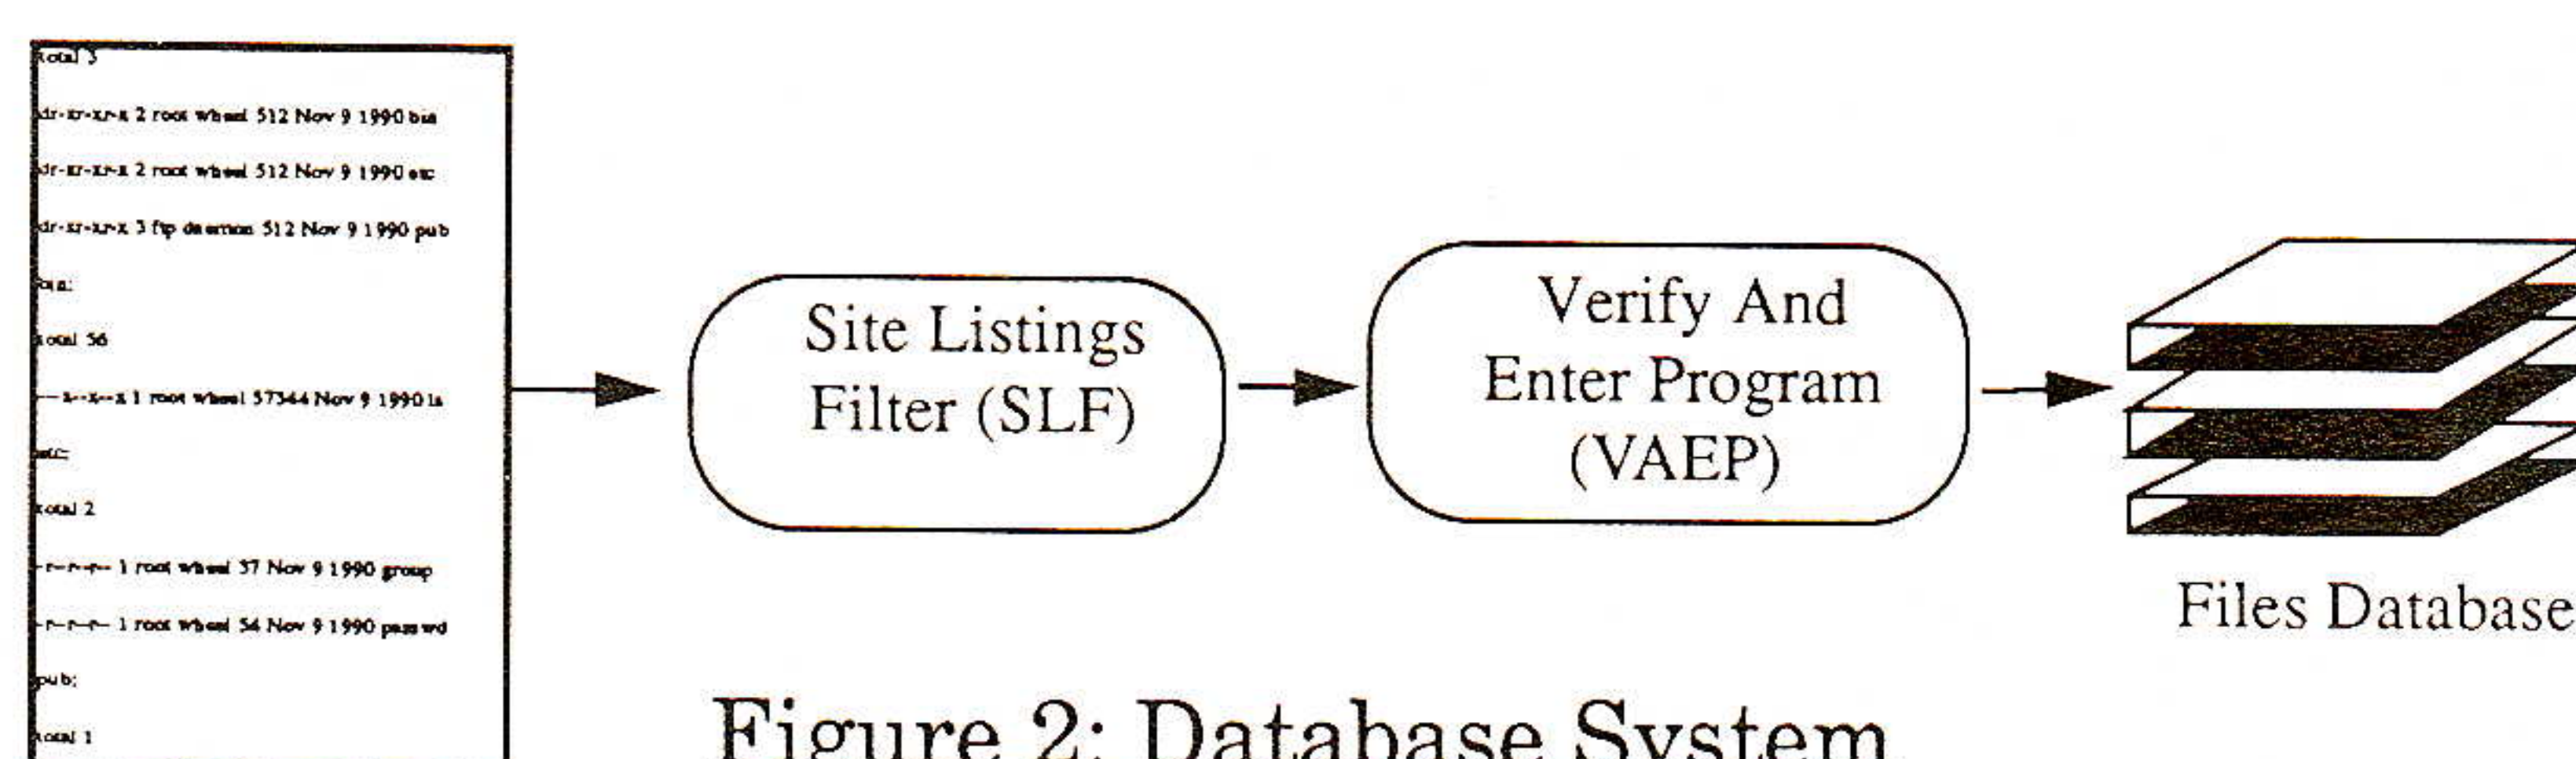


Figure 2: Database System

The first program is the *Site Listings Filter* (SLF) which removes erroneous information (such as error messages generated by the *ls(1)* command) from a raw site listing. The clean site listing is then passed to the *Verify And Enter Program* (VAEP). The VAEP parses the clean listing file, rebuilding the directory hierarchy of the site in memory to verify the listing's consistency.

Once the directory structure is verified correct, the VAEP scans the tree, inserting the information into the *archie* Files Database. If the verification step fails, the update aborts and the operator is notified.

continued on next page

The *archie* System (continued)

There is no attempt made to allow partial insertions or updates to a site listing. Given the dedicated format used for the Files Database and the problems that would have been encountered in maintaining consistency while processing updates on an active database, it was decided that writing and debugging the needed code was not worth the effort. Instead, a site's entries are all deleted using a separate delete program before running the VAEP for that site. Although this approach makes the VAEP (and thus site updates) computationally expensive to perform, it also made implementation easier and allowed us to get the new system up and running that much faster.

As with many other parts of *archie*, the lack of resources (especially time for development) was a major factor influencing our design decisions. In this case we have traded runtime resources needed to perform deletion and re-insertion for ease of implementation and maintenance. In practice it is a decision that only *archie* site administrators have regretted. A new database format has been designed for the next version of *archie* and these design decisions are being re-examined.

User Access Component

The *User Access Component* (UAC) allows individual Internet users to access and query the various *archie* databases using a variety of access methods or channels. These include *telnet*(1), electronic mail, or through the *Prospero* file system protocol. Work is also underway to make the *archie* databases available directly through the *Wide Area Information Server* (WAIS) and *WorldWideWeb* (WWW) systems. Collaboration with other projects is welcomed.

Telnet Command Interpreter

The first user interface channel to the *archie* databases was provided by a simple command line interpreter, a version of which runs on each *archie* server. Access to this *Telnet Command Interpreter* (TCI) is through the *telnet*(1) command, which is assumed to be available on most Internet connected sites.

The TCI provides full functionality to the *archie* system using a simple (and relatively primitive) interface. Users can specify searches in either of the available databases or access information about each site. There is also access to an e-mail interface to have either search results or the *archie* manual page sent back via e-mail. Users can also access online help, list available *archie* servers or manipulate a number of variables that are used to control operation.

In operation, the TCI has proved to be a serious resource drain under high load. Each *archie* telnet session requires a copy of the command interpreter to be launched, and each instantiation requires a large number of open files to access the various components of the databases, along with significant amounts of other machine resources (such as core memory, swap space, etc.).

Each instantiation of the TCI accesses the single copy of the *archie* Files Database. This database is large (currently over 110 Megabytes) and thus the current version of the access routines maps the appropriate files into memory using the SunOS *mmap*(2) call to improve performance (a corresponding functionality is used in the latest servers, which operate on the IBM RS-6000 class machines). This allows reasonable response time, but means that an *archie* server will benefit from all the RAM its owners can provide.

As the popularity of *archie* has grown, it was not uncommon to see over 40 simultaneous *telnet* sessions, at which point the server would become almost unusable.

To address these problems, a limit has been placed on the number of simultaneous *archie* login sessions at the pilot Montreal server. This was done after a client-server access model became available with the arrival of the *Prospero* user interface (see below). Since this was done, total system throughput (as measured in the number of file database searches per day) has in fact gone up, since the *Prospero* interface is far less resource intensive.

There are plans to rewrite the TCI so that it uses a client-server access model. The idea is to have the current TCI generate queries and send them to the *Prospero* interface using the UDP-based *Prospero* protocol. This would address the problem of machine resources (only one set of open database file pointers would be needed within the *Prospero* interface, for example). It would also allow the *telnet* interface to access other resource providers (such as an *archie* WAIS interface) as they are developed.

E-mail Interface

Historically, the *E-mail Interface Server* (EIS) was the second developed. Users can send in queries to the *archie* databases via e-mail to the EIS, which performs the specified search and returns the results in an e-mail message. The EIS is based in concept upon the KISS mail server package, available from a number of archive sites.

Functionality of the e-mail interface has always lagged behind that of the TCI. For example, the e-mail interface does not currently support the ability (present in the TCI) to set variables to control system operation. Rationalizing the various user interface channels to permit a consistent view and full functionality through all mechanisms is yet another item that we have appended to the list of things to be addressed as time permits.

Although the *archie* system itself is not capable of performing an anonymous FTP transfer for users of the e-mail interface, there is a system operated by Digital Equipment Corporation that will perform such fetches via e-mail. Details on both the *archie* e-mail interface and the DEC e-mail anonymous FTP services can be obtained by sending an e-mail message to `archie@archie.mcgill.ca` with the word "help" in either the subject or message body.

Prospero Interface

In early June, 1991 we entered into a successful collaboration with Clifford Neuman of ISI, when he ported his *Prospero* file server to the *archie* system, giving us the *archie Prospero Interface Server* (PIS). The PIS allows users of the *Prospero* system to access the *archie* files and Whatis databases through the *Prospero* server without the need to log onto the *archie* server directly.

The *Prospero* system uses a UDP-based protocol that is far less resource intensive than the *telnet*(1) client. The server architecture also allows the use of sophisticated query scheduling algorithms for selecting queries to be performed. This is useful because the different types of available searches have widely varying impact on system performance. For example, exact match searches can be performed in $O(1)$ time, while full regular expression matches take $O(n)$, where n is the number of unique strings in the database. In operation, the PIS query scheduler will give preference to exact match requests to maximize throughput.

The existence of the *Prospero archie* server has spurred the development of a number of stand-alone *archie* client programs based upon the *Prospero* protocol. These now include a stand-alone command line version (one that runs on the user's machine, not the *archie* server), an X Windows version, as well as others.

continued on next page

The *archie* System (continued)

Future work on the user interface

These programs are available to users from a number of Internet archives, including the anonymous archive on `archie.mcgill.ca`.

Clifford Neuman continues to work with us on improving the *Prospero archie* interface and we have elected to standardize our current client-server efforts for accessing the Files Database on this channel. Steps must still be taken to extend the full functionality of the TCI interface to the *Prospero* server and this is planned.

The work on the *Prospero* interface has been followed by recent efforts by Brewster Kahle of Thinking Machines Inc., who has been investigating the possibility of making the *archie* databases available via the WAIS system. Initially, the information in the *archie* Files Database has simply been reformatted into a single huge text file and then indexed using a WAIS server.

Although this does make the information available, it is very resource intensive and presents problems with updating. Whenever the database is modified the index must be regenerated. Thus each *archie* database update is potentially a very CPU-intensive operation. Adequately addressing this problem is the subject of on-going research.

In the long term we would like to create a WAIS server for the *archie* system to permit a complete WAIS interface to the *archie* databases. We also plan to deploy a number of additional databases and are investigating the possibility of using WAIS as our search and retrieval engine for accessing these. Most of our planned offerings will take the form of large textual databases which make them ideal candidates for the WAIS system.

Using *archie*

Users with direct Internet connectivity can try out an interactive *archie* server using the basic *telnet* command (available at most sites). To use, *telnet* to the host `archie.mcgill.ca` [132.206.2.3] and login as user "archie" (there is no password needed). A banner message giving latest developments and information on the *archie* project will be displayed and then the command prompt will appear. First-time users should try the "help" command to get started.

Users with only e-mail connectivity to the Internet should send a message to `archie@archie.mcgill.ca`, with the single word "help" in either the subject line or body of the message. You should receive back an e-mail message explaining how to use the e-mail *archie* server, along with details of an e-mail-based FTP server operated by Digital Equipment Corporation that will perform FTP transfers through e-mail requests.

Demo *archie* clients are stored on `archie.mcgill.ca` in the sub-directory `archie/clients` and may be obtained using anonymous FTP. There are several such clients and others are currently being tested. Additional work is planned in this area in the coming months and details will be announced in the *archie* banner message displayed on login.

Documentation for the *archie* system is still limited, but what there is is also available for anonymous FTP from the same host under the directory `archie/pub`. This includes a UNIX-style manual page, as well as an on-line copy of this article.

The *archie* project continues to grow in part because of the feedback and response from users. Suggestions for improvements and additional features are especially welcome. Please let us know what you think!

Contacting the archie people

Please send comments, suggestions and bug reports to `archie-group@archie.mcgill.ca`. This address reaches the implementors of *archie*. There is also the `archie-people@archie.mcgill.ca` mailing list. This list is for people interested in developments and progress of the *archie* project and is open to all who wish to subscribe.

Surface mail address:

UNIX Support Group
Computing Centre
McGill University
Room 200, Burnside Hall
805 Sherbrooke Street West
Montreal, Quebec
CANADA H3A 2K6
Phone: 514-398-3709
E-mail: `peterd@cc.mcgill.ca`

References

- [1] Kahle, Brewster, "An Information System for Corporate Users: Wide Area Information Servers," *ConneXions*, Volume 5, No. 11, November, 1991.
- [2] Franklin Davis et al., "WAIS Interface Protocol Prototype Functional Specification," Thinking Machines, Corp. Available from Franklin Davis (`fad@think.com`) or Brewster Kahle (`brewster@think.com`).
- [3] Barron, Billy, "Another use of the Internet: Libraries Online Catalogs," *ConneXions*, Volume 5, No. 7, July 1991.
- [4] Quarterman, John, "Networks: From Technology to Community," *ConneXions*, Volume 5, No. 7, July 1991.
- [5] Schwartz, Michael F., "Resource Discovery and Related Research at the University of Colorado," *ConneXions*, Volume 5, No. 5, May 1991.
- [6] Malamud, Carl, "Is Resource Discovery Hacking?," *ConneXions*, Volume 5, No. 12, December 1991.
- [7] Huston, Geoff, "Profile: AARNet—The Australian Academic and Research Network," *ConneXions*, Volume 4, No. 3, March 1990.

PETER DEUTSCH received his B.Sc. in Mathematics & Computer Science from McGill University in 1985. He is currently working on his M.Sc. in Computer Science at the same institution, expecting to finish this year. Prior to this, he was employed by McGill University as the Systems Manager for the School of Computer Science during a period in which his group installed the first Internet backbone connection into Montreal, installed one of the first USENET feeds, installed the first Montreal BITNET to Internet gateway, and generally brought the Internet to Quebec. He went on to set up the UNIX Support Group at the McGill University Computing Centre.

ALAN EMTAGE received his B.Sc. in Mathematics and Computer Science in 1987 and M.Sc. (Applied) in Computer Science in 1991, also from McGill University. He worked as a Systems Programmer for McGill from 1986 until 1991, first for the School of Computer Science and then for the UNIX Support Group of the Computing Centre.

The two authors are deeply interested in the various issues involved in developing and providing user services in an Internet environment and are the principal architects of the *archie* system, which grew out of their studies. They have recently formed a company to develop and market Internet information discovery and retrieval tools. They are also co-chairs of the Internet Engineering Task Force working group on *Internet Anonymous FTP Archives* (IAFA-WG), a group which is investigating ways to improve archiving and information delivery services on the Internet.

ISDN as an Open Platform for Telecommunications Infrastructure

by Mitchell Kapor, The Electronic Frontier Foundation

Background

This overview is a summary of testimony presented by the *Electronic Frontier Foundation* (EFF) to the House Energy and Commerce Committee, Subcommittee on Telecommunications and Finance in hearings regarding Telecommunications Infrastructure Legislation and Proposals, October 24, 1991. The testimony was prepared by Mitchell Kapor in consultation with Jerry Berman, Director of the ACLU Information Technology Project and Danny Weitzner. Many people in the computer and networking community also contributed valuable comments and suggestions.

The infrastructure challenge

Until now the telecommunications policy debate has largely been framed as a struggle among entrenched commercial interests over who will control and dominate markets such as information services, manufacturing, and long distance service. With this proposal we hope to help to refocus the debate by defining public goals and specific steps that can be taken to achieve them. Public policy should be guided by an overarching social vision of what we call the *National Public Network* (NPN), a vibrant web of information links to serve as the main channels for commerce learning, education, politics, social welfare, and entertainment in the future. This network will include the voice telephone service that we are already so familiar with, along with video images, sound, and hybrid forms of communication.

In the view of EFF we need more than just safeguards, entry level tests or new telephone company investment in information services and fiber optics. In order to ensure a level playing field, encourage diversity, and safeguard the freedom of users, we must build an open telecommunications platform according to the following principles:

- Establish an open platform for information services by speedy deployment of "Personal ISDN" nation-wide;
- Ensure competition in local exchange services;
- Promote First Amendment free expression by reaffirming the principles of common carriage;
- Foster innovations that make networks and information services easy to use;
- Protect personal privacy; and
- Preserve and enhance socially equitable access to communications media.

Policy recommendations

I. Create an Open Platform for Innovation in Information Services by Speedily Deploying a Nation-wide, Affordable "Personal ISDN": To achieve the information diversity currently available in print and broadcast media in the new digital forum, we must guarantee widespread accessibility to a platform of basic services necessary for creating information services of all kinds. Such a platform offers the dual benefit of helping to create a level playing field for competition in the information services market, and stimulating the development of new services beneficial to consumers.

Some suggest that the technology necessary to offer such a platform is far off and would require billions of dollars of investment in fiber optics. Actually, we have a platform that meets these criteria within our reach right now.

ISDN Personal ISDN (*Integrated Services Digital Network*) [2, 3] could make voice, data, video, high-speed fax, and multimedia services available *today* to telephone subscribers all around the country. ISDN as a key information services technology is well-known in the communications industry, but its potential as a universal platform is not properly appreciated, nor has it been properly priced and positioned by the RBOCs as a basic service for everyone, including consumers and small businesses.

Characteristics The desktop personal computer represented a revolutionary platform for innovation of the 1980's because it was designed according to the principle of *open architecture*, allowing numerous hardware and software entrepreneurs to enter the computer industry. To bring the benefits of the information age to the American public in the 1990's, we need to build an open, ubiquitous digital communications platform for information services. Personal ISDN can enable the citizen's access into the Information Age because it has these key characteristics:

- *Critical Mass of Features:* Existing ISDN standards, once fully implemented, offer switched, high-speed, error-free data communications which can deliver a variety of advanced information services. Many of the capabilities once thought to be possible only on an all-fiber network, such as interactive full-motion video can be achieved to a significant degree over Personal ISDN. This is due to continuing revolutions in compression technology which makes it possible to use copper wire-based ISDN to carry video signals to their destination, at which point they are uncompressed through use of increasingly inexpensive processors, which are built-in to computers, televisions, and other consumer electronic equipment.
- *Ubiquity:* To create a market for information services, everyone must be able to reach the platform. We must build the new public network by making it easy for people to connect to it with a few simple decisions. Again, an analogy to the personal computer market is helpful. Minicomputers and mainframes were marketed to companies. Microcomputers (PCs) were marketed to individuals. Personal ISDN—which can be provided over the existing copper plant that comprises today's public switched network—can reach into every home and every small business without laying a single mile of fiber optic cable. Telephone company data indicates that over the next three years the majority of all US central office switches will be upgraded to the requisite digital capability.
- *Affordability:* Platform services, even if they are ubiquitous, are useless unless they are also affordable to American consumers. Just as the voice telephone network would be of little value if only a small fraction of the country could afford to have a telephone in their home, a national information platform will only achieve its full potential when a large majority of Americans can buy access to it. All available information indicates that ISDN can be priced as a basic service. The cost of carrying a digital ISDN call from the customer to the local switch is just the same as an analog voice call in the digital switching regime that ISDN presupposes. There are some fixed investment costs still to be incurred to upgrade the nation's central office switches in order to handle ISDN traffic, but commitments to these investments have already largely been made.

ISDN as an Open Platform (*continued*)

What is needed is to raise the floor by creating a new standard, minimum platform for information exchange. ISDN must be repositioned as a basic service, available to consumers and small businesses. This service can be the testbed for a whole new generation of information services which could benefit the American public and level the competitive playing field.

Competition

II. Ensure Competition in Local Exchange Services: Many consumer and industry groups are concerned that as the MFJ [1] restrictions are lifted, the RBOCs will come to dominate the design of the emerging National Public Network, shaping it more to accommodate their business goals than the public interest. The bottleneck that RBOCs have on local exchange services critical to information providers can be minimized by unbundling these services and allowing non-BOC providers to offer them in competition with BOC local exchange companies.

Some suggest that an entry level test is necessary to guarantee that alternative infrastructure is developed for information services delivery. Alternative pathways are a useful and necessary part of our telecommunications infrastructure, but we should not rely on them alone to level the information services playing field. First and foremost we must find ways to open up the existing public switched network to competition at all levels. Competition will promote innovation in the services on which information providers rely, and help guarantee equal access to all local exchange facilities.

The post-divestiture phone system offers us a valuable lesson: a telecommunications network can be managed effectively by separate companies—even including bitter opponents like AT&T and MCI—as long as they can connect equitably and seamlessly from the user's standpoint. Together with the open platform offered by a Personal ISDN, unbundling and expanded competition is a key to ensuring equitable access to local exchange services needed for information service delivery.

Common Carriage

III. Promote First Amendment Free Expression by Affirming the Principles of Common Carriage: In a society which relies more and more on electronic communications media as its primary conduit for expression, full support for First Amendment values requires extension of the common carrier principle to all of these new media. Common carriers are companies which provide conduit services for the general public. The common carriers' duties have evolved over hundreds of years in the common law and later statutory provisions. The rules governing their conduct can be roughly distilled in a few basic principles. Common carriers have a duty to provide services in a non-discriminatory manner at a fair price, interconnect with other carriers, and provide adequate services. The communications carriers who make up the critical elements of the public switched network—local exchange companies and inter-exchange companies—should be subject to comprehensive common carriage duties as described above. All communications carriers, however, are not necessarily common carriers.

Unlike arrangements found in many countries, our communications infrastructure is owned by private corporations instead of by the government. Therefore, a legislatively imposed expanded duty of common carriage on public switched telephone carriers is necessary to protect free expression effectively.

A telecommunications provider under a common carrier obligation would have to carry any legal message regardless of its content whether it is voice, data, images, or sound. For example, if full common-carrier protections were in place for all of the conduit services offered by the phone company, the terminations of "controversial" 900 services such as political fundraising would not be allowed, just as the phone company is now prohibited by the Communications Act from discriminating in the provision of basic telephone services.

Simplicity

IV. Make the Network Simple to Use: One of the great virtues of today's public switched telephone network, from a user's perspective, is that it operates according to patterns and principles that are now intuitively obvious to almost everyone. As this network grows beyond just voice services, information services that become part of this network should reflect this same ease-of-use and accessibility. The development of such standards and patterns for information services is vital, not just because it helps make the network easier to use, but also because it ensures an open platform for information providers. However, standards development will be *ad hoc* and even chaotic at first. Numerous standards may be tried and found inadequate by users before a mature set of standards emerges. Congress and government regulatory bodies may need to set out the ground rules for standards planning in order to ensure that all interested parties have an equal voice, and the resulting standards should be closely analyzed to make sure that they reflect public needs. But, direct government involvement in the process should be as limited as possible.

Privacy

V. Protect Personal Privacy: As the NPN develops, there are threats to both communications privacy and information privacy. First, electronic communications meant to be private can be intercepted without the consent or even knowledge of the communicating parties. The privacy of telephone conversations and electronic mail is already protected by the *Electronic Communications Privacy Act*. However, communications in other media, such as cellular phone conversations, can be intercepted using readily available technology by private third parties without the knowledge or consent of the conversants. Second, as the public switched telephone network is used for an increasing variety of transactions, it will hold more personal information about consumers. We need to give citizens greater control over information collected, stored, and disseminated by telephone companies and information providers. As the public outcry over Caller ID demonstrates, citizens want and deserve to have adequate notice about what information is being collected and disseminated by communications firms and must be able to exercise informed consent before information collected for one purpose can be used for any other purpose.

Equitable Access

VI. Preserve and Enhance Socially Equitable Access to Communications Media: The principle of equitable access to basic services is an integral part of nation's public switched telephone network. We must ensure that all Americans have access to the growing information services market. Some paint a vision of the future in which all citizens have access to education services such as distance learning or on-line health care services. Neither market competition nor lifting restrictions on telephone companies alone will deliver these services. It is time for those who propose serving the "information have-nots" to admit that equity cannot be achieved except by legislative mandate and public funding.

ISDN as an Open Platform (*continued*)

Conclusion

The chance to influence the shape of a new medium usually arrives when it is too late: when the medium is frozen in place. Today, because we are at the cross-roads of telecommunications policy, and because of the unusual awareness people have of its possibilities, there is a rare opportunity to shape this new medium in the public interest, without sacrificing diversity or financial return.

More information

For a copy of the complete testimony on which this overview is based or for more information please contact:

Mitchell Kapor, President
Electronic Frontier Foundation
155 Second St.
Cambridge, MA 02141
617-864-0665
mkapor@eff.org

or:

Daniel J. Weitzner
EFF Washington Office
666 Pennsylvania Ave, SE
Suite 303
Washington, DC 20003
212-544-9237
dweitzner@eff.org

EFF

The Electronic Frontier Foundation (EFF) is a public interest organization established in 1990 to educate the public about the democratic potential of new computer and communications technologies. EFF works to develop, and seeks to implement, public policies to maximize freedom, competitiveness, and civil liberty in the electronic social environments being created by these new technologies.

Documents

"Open Platform Overview": This is the document you are now reading. It summarizes our policy recommendations for the creation of a ubiquitous, affordable, open telecommunications platform based on ISDN. A slightly different version was printed in *EFFector* 2.01. Additional copies may be obtained via electronic mail. Send a message to archive-server@eff.org, any subject, with body: send documents open-platform-overview. The document is also available via anonymous FTP from [eff.org/pub/docs/open-platform-overview](ftp://eff.org/pub/docs/open-platform-overview).

"Testimony of Mitchell Kapor Before the House Subcommittee on Telecommunications and Finance Regarding Telecommunications Infrastructure Legislation and Proposals": This is the complete testimony presented to Congress, which is the full text from which the overview was prepared. The testimony can be obtained via electronic mail by sending a message to archive-server@eff.org, with any subject, with body: send documents open-platform-testimony, or via anonymous FTP from [eff.org/pub/docs/open-platform-testimony](ftp://eff.org/pub/docs/open-platform-testimony).

EFFector Online: This is the regular newsletter of the Electronic Frontier Foundation. We will continue to report progress on the Open Platform initiative here. The newsletter is available via electronic mail. Send e-mail to eff-news-request@eff.org requesting to be put on the mailing list. Or read it on USENET in the group <comp.org.eff.news>.

IBT mailing list: The IBT (*Internet Brain Trust*) moderated mailing list is being organized as a forum for discussion on the Open Platform. To join the list, please send mail to `ibt-request@eff.org`. The IBT archive will be available via FTP from `eff.org:/pub/ibt`.

General information about the EFF, including membership information can be obtained via electronic mail or FTP. Send mail to `archive-server@eff.org`, any subject, with body: send EFF EFF.about or get the file `eff.org:/pub/EFF/EFF.about` via anonymous FTP.

References

- [1] H. H. Greene, United States of America, Plaintiff, v. Western Electric Company, Inc., et al., Defendants. Civil Action No. 82-0192, U.S. District Court, District of Columbia, March 1988, "Triennial review of Modified Final Judgement divesting AT&T of the Regional Bell Operating Companies."
- [2] Blackshaw, R., "Components of OSI: Integrated Services Digital Networks (ISDN)," *ConneXions*, Volume 3, No. 4, April 1989.
- [3] Leifer, D., "ISDN: Why use it?," *ConneXions*, Volume 4, No. 10, October 1990.

[Ed.: See also "Call for Nominations: The Electronic Frontier Foundation's First Annual Pioneer Awards," on page 21].

MITCHELL KAPOR is the co-founder and President of the Electronic Frontier Foundation, an organization concerned with educating the public and policy makers about computer-based communications media and information networks. Mr. Kapor is active as a speaker and writer about civil liberties and electronic media, the national information infrastructure, intellectual property, and software design. Previously, Mr. Kapor founded Lotus Development Corporation and served as its Chief Executive Officer, President, and Chairman. He is the designer of *Lotus 1-2-3* and many other software applications. Currently, he serves as Chairman of ON Technology, Inc. of Cambridge, Massachusetts, a developer of local-area network applications for collaborative computing. He received his B.A. from Yale College in 1971 in a self-designed major combining psychology, linguistics, and computer science. He has a Master's degree in psychology and has studied management at MIT's Sloan School. Mr. Kapor has served on several government advisory panels and boards in the areas of computer science, intellectual property rights for software, and information infrastructure.

Errata

ConneXions is by no means immune from the typo-disease. In our December 1991 (Volume 5, No. 12) issue we sincerely apologize for the two following errors:

- Joyce Reynolds works for the USC *Information Sciences Institute* (ISI), not for SRI as we erroneously reported. (page 11).
- Mitch Kapor represents the *Electronic Frontier Foundation*, not "...Freedom Foundation." (page 11).

...and of course, on page 1, I meant to say "...have a good holiday," my standard excuse is that English is only my second language :-).

—Ole

NREN Bill Signed into Law

by Mike Roberts, EDUCOM

Introduction

After a lengthy and tortuous legislative history, the *High Performance Computing Act* of 1991 was signed into law by President Bush on December 9th, 1991. The bill provides a mandate for federal support of a variety of computing and communication activities and authorizes expenditure of more than three billion dollars over the next five years for existing and new programs. This is equal to about 1% of total planned federal R&D expenditures over the same period. The amount allocated to the NREN, for both research and production expenses, is approximately \$100 million a year over the five years. Appropriations for agencies participating in the HPCC program in FY92 were recently completed with requested funding levels essentially intact.

In his remarks at the signing ceremony, the President said, "The development of high performance computing and communications technology offers the potential to transform radically the way in which all Americans will work, learn, and communicate in the future. It holds the promise of changing society as much as the other great inventions of the 20th Century, including the telephone, air travel and radio and TV."

Two approaches

For the last two years, Republicans in the Administration and Democrats in the Congress have been pursuing separate but parallel approaches to high performance computing and communications (HPCC). For the Administration, the key players have been Science Advisor Allan Bromley and Budget Director Richard Darman. In the Congress, Senator Al Gore and Representative George Brown, chairmen of the respective Science Committees, teamed up to push the bill through in 1991 after a last minute series of glitches ended in failure last year. Despite moments of partisan politics, a cast of dozens of hard working advocates for the NREN, in and out of government, persevered to a successful conclusion.

Originally conceived as a high tech replacement for the aging and overloaded ARPANET, plans for the NREN have undergone successive waves of program redefinition and expansion over the past five years. For instance, section 102 of the bill states, "The Network (i.e., NREN) shall provide for the linkage of research institutions and educational institutions, government, and industry in every state."

Benefits

As the benefits of computer-based networking have grown more apparent in recent years, pressure for greater access to the NREN has increased. The access requirement, as passed by the Congress, says, "Federal agencies and departments shall work with private network service providers, State and local agencies, libraries, educational institutions and organizations, and others, as appropriate, in order to ensure that the researchers, educators and students have access, as appropriate, to the Network. The Network is to provide users with appropriate access to high performance computing systems, electronic information resources, other research facilities, and libraries. The Network shall provide access, to the extent practicable, to electronic information resources maintained by libraries, research facilities, publishers, and affiliated organizations."

In another change from earlier plans, the bill makes provision for the uses of the network in addition to its creation. "The Director (of the White House Office of Science and Technology Policy [OSTP]) shall assist the President in coordinating the activities of appropriate agencies and departments to promote the development of information services that could be provided over the Network."

These services may include the provision of directories of the users and services on computer networks, data bases of unclassified Federal scientific data, training of users of data bases and computer networks, access to commercial information services for users of the Network, and technology to support computer-based collaboration that allows researchers and educators around the Nation to share information and instrumentation."

Characteristics

The enabling legislation explicitly recognizes that many of the goals of the NREN cannot be realized with today's technology. In addition to the requirement that the network demonstrate gigabit transmission speeds by 1996, it sets out ten desired characteristics of the network as a guide for its development and evolution.

"The Network shall:

- (1) be developed and deployed with the (assistance of) the computer, telecommunications, and information industries;
- (2) be designed, developed, and operated in collaboration with potential users in government, industry, and research institutions and educational institutions;
- (3) be designed, developed, and operated in a manner which fosters and maintains competition and private sector investment in high speed data networking within the telecommunications industry;
- (4) be designed, developed, and operated in a manner which promotes research and development leading to development of commercial data communications and telecommunications standards, whose development will encourage the establishment of privately operated high-speed commercial networks;
- (5) be designed and operated so as to ensure the continued application of laws that provide network and information resources security measures, including those that protect copyright and other intellectual property rights, and those that control access to data bases and protect national security;
- (6) have accounting mechanisms which allow users or groups of users to be charged for their usage of copyrighted materials available over the Network and, where appropriate and technically feasible, for their usage of the Network;
- (7) ensure the interoperability of Federal and non-Federal computer networks, to the extent appropriate, in a way that allows autonomy for each component network;
- (8) be developed by purchasing standard commercial transmission and network services from vendors whenever feasible, and by contracting for customized services when not feasible, in order to minimize Federal investment in network hardware;
- (9) support research and development of networking software and hardware; and
- (10) serve as a testbed for further research and development of high capacity and high-speed computing networks and demonstrate how advanced computers, high-capacity and high-speed computing networks, and data bases can improve the national information infrastructure."

NREN Bill Signed into Law (*continued*)

Suspensions

The Congress has some lingering suspicions about the depth of the Republican Administration's commitment to the NREN, and has provided in the new law for a series of studies and reports. Within a year of enactment, the Director of OSTP, Dr. Bromley, is to report to the Congress on issues such as combining commercial and non-commercial services, funding sources, operational structure, security policies, and protection of copyrighted material distributed on the Network. It is expected that individuals and organizations from throughout the Internet community will be asked to contribute to these studies and accompanying recommendations.

Management

Over the multi-year legislative history of the NREN bill, unanimity emerged from industry, academic, and agency testimony on the strong need for a high-performance national network. The final bill is remarkably faithful to the basic principles laid out in 1986 and 1987 despite the broader constituencies and responsibilities picked up along the way. The chief failings of the measure are in the areas of governance and management, and reflect both agency politics and genuine uncertainty in the community on the best way of building and operating such an ambitious undertaking. The federal agencies with the largest stakes in the NREN—NSF, Energy, Defense, and NASA—were unwilling to have one amongst them singled out for a lead agency management role. They and their Congressional committee supporters were even unwilling to have the Director of OSTP tasked with management of the NREN. The final language turns the responsibility over to the President, which provides the maximum maneuvering room at the Cabinet Secretary level for behind the scenes horsetrading on assignment of network management responsibility.

Of perhaps greater consequence, the bill makes no provision for representative governance of the NREN, a network which will require vastly larger investments outside of the federal government than within it. The tolerance within the networking community for the excessive federal bias of the program, largely the product of the quietly effective way in which federal networking executives have supported Internet growth in the last several years, may be sorely tested if reasonable mechanisms for shared governance of this unique enterprise are not forthcoming.

References

- [1] Roberts, M., "The NREN Takes Shape," *ConneXions*, Volume 5, No. 6, June 1991.
- [2] Dern, D., "The ARPANET is Twenty," *ConneXions*, Volume 3, No. 10, October 1989.
- [3] NNSC Staff, "Profile: NSFNET," *ConneXions*, Volume 1, No. 2, June 1987.
- [4] Braun, H-W., "The new NSFNET backbone network," *ConneXions*, Volume 2, No. 12, December 1988.

MIKE ROBERTS is Vice President for Networking at EDUCOM, a 600 member association of colleges and universities with common interests in information technology. The EDUCOM Networking and Telecommunications Task Force, a group of sixty universities and corporations, of which he is the staff director, has been active in planning and advocacy for the NREN. His Internet address is roberts@educom.edu.

Book Reviews

There are a plethora of books out on OSI. This month we present micro-reviews of two more.

Implementing OSI Networks, Gerald D. Cole, John Wiley & Sons, Inc., ISBN 0-471-51060-2 (paper), 350 pp.

Broad overview

This book is a broad overview of OSI, and is excellent for that purpose. However, it has little to do with OSI implementation issues—either in coding or deployment.

From the perspective of a technical overview, the book covers the basics of each layer, along with the major applications. In addition, it has a brief section on OSI profiles and testing. The final chapter (“Evolving OSI Developments”) is a stream-of-consciousness discussion on OSI-related topics such as standards-making and transition. This last chapter is where the book suffers most. The author introduces many interesting issues, but the treatment of each is entirely too brief. It is enough to whet the appetite of the reader, and perhaps slightly misinform, due to not covering all the angles. However, the terse coverage is probably adequate given the book’s scope.

Big Picture

Implementing OSI Networks makes a fine introduction to OSI technology. It’s not big on details, but you do get the Big Picture.

Open Systems Interconnection Handbook, Gary R. McClain (editor), Intertext Publications, ISBN 0-07-044969-4 (cloth), 412 pp.

This book is a collection of articles, mostly on OSI topics. There are four sections: Standards and Protocols, Networks and Architectures, X.400 and X.500, and Implementation Considerations.

Hype

Unfortunately, most of the authors seem to have captured the hype of OSI without having an understanding as to what makes OSI good or bad. There are some exceptions: for example, Ashar Aziz of Sun Microsystems contributed one of the finest descriptions of the OSI network layer that I’ve ever read. It’s a pity that most of the authors didn’t contribute articles of similar quality. The editor gets low marks for a rather confusing organization. For example, articles on “Successful OSI Migration Strategies” and the OSF/DCE somehow got bundled into the X.400/X.500 section. What the DCE has to do with X.400/X.500—or OSI in general—is anyone’s guess. Further, I wonder if the authors of the article on migration strategies understand that their title implies that their test cases have successfully migrated from one protocol suite to OSI and then back again, probably on a seasonal basis.

Shameless

The book also contains some shameless advertising—each article ends with information about the author(s). Some of these articles then go on to extol the virtues of the authors’ employers. *Puh-leeze!*

I generally refuse to review a book if I can’t find anything good to say about it. The only reason that I agreed to review *Open Systems Interconnection Handbook* for the industry newsletter* *ConneXions* was because of a few articles like the one by Aziz. Other than that ...

—Marshall Rose

*Ed.: We like to call it “monthly technical journal,” but Dr. Rose insists on this label!

Is Resource Discovery Hacking?—A reader responds

Carl,

Thank you for a very provocative article [*ConneXions*, Volume 5, No. 11, November 1991]. I must say, however, that I disagree with your conclusions. You seem to be saying that “resource discovery upsets a lot of system administrators who consider it trespassing, but it’s really not, because, um, well, uh, because it’s not.”

I know that all kinds of analogies have been offered, and endless discussions, but I still feel I have to add my piece. We agree that some parts of the Internet are private, and some are public, but we seem to disagree as to where to draw the line. You draw the line a little bit inside my computers, I’d draw the line at the router that connects me to our regional network.

My analogy is that our computers are like my real estate property. We both agree that you can’t come into my house uninvited, but you think walking on my grass is OK. I don’t, simply for the reason that it’s my grass, and I make the rules for it. Yes, I know that occasionally cutting across my yard won’t hurt the grass, and in fact there may be benefits to the community at large if I allow it (people get to work quicker), but that doesn’t change a thing. It’s *my* grass, and I can say who walks on it. If you ask, I might very well say “sure,” but you must ask.

The same thing for my computer systems. I don’t want strangers (which includes you and Schwartz) using a single cycle of my CPUs. That may seem irrational, since you can totally dominate my system with anonymous FTP requests, but that’s different, because I advertise that anyone can use my system for anonymous FTP. I *didn’t* say you could use it for any kind of research, no matter how few CPU cycles you consume. You may feel that your research is benign, but I might not. They’re my CPU cycles, and I’d like the right to make the determination.

I think the real place to address this is with the contracts that users have in place with the regional networks. We pay xxxx dollars a year for network connectivity. Somewhere in that contract it should say: “by the way, by being on the Internet you agree to contribute xx% of your machine cycles to Internet research.” I feel users should have the option of saying, “no, I’m not interested in that,” and the list of non-participants be made available to researchers.

—Michael H. Morse, National Science Foundation

The author responds:

While you may own your house, you give up absolute rights as part of the price of being in a community. I view your border system as being yours, but also serving a community purpose. Instead of treating your border system as your front lawn, you should view part of it as being a sidewalk.

It is important that we do not forget that the Internet is much more than just a transmission mechanism. It is a *community* in which we all participate. We must occasionally subjugate absolute control over property in the broader interests of the community. I believe that it is this philosophy that led the Internet Activities Board to decide that activities such as those of Mike Schwartz are not only allowable, but a desirable part of the Internet.

—Carl Malamud

Call for Nominations:

The Electronic Frontier Foundation's First Annual Pioneer Awards

Background

In every field of human endeavor, there are those dedicated to expanding knowledge, freedom, efficiency and utility. Along the electronic frontier, this is especially true. To recognize this, the Electronic Frontier Foundation has established the *Pioneer Awards*. The first annual Pioneer Awards will be given at the Second Annual Computers, Freedom, and Privacy Conference in Washington, D.C. in March of 1992.

All valid nominations will be reviewed by a panel of outside judges chosen for their knowledge of computer-based communications and the technical, legal, and social issues involved in networking.

Guidelines

There are no specific categories for the Pioneer Awards, but the following guidelines apply:

- The nominees must have made a substantial contribution to the health, growth, accessibility, or freedom of computer-based communications.
- The contribution may be technical, social, economic, or cultural.
- Nominations may be of individuals, systems, or organizations in the private or public sectors.
- Nominations are open to all, and you may nominate more than one recipient. You may nominate yourself or your organization.
- All nominations, to be valid, must contain your reasons, however brief, on why you are nominating the individual or organization, along with a means of contacting the nominee, and your own contact number. No anonymous nominations will be allowed.
- Every person or organization, with the single exception of EFF staff members, are eligible for Pioneer Awards.

Sending nominations

You may nominate as many as you wish, but please use one submission per nomination. You may send us the nominations at:

Pioneer Awards
EFF
155 Second Street
Cambridge, MA 02141
Fax: 617-864-0866
E-mail: pioneer@eff.org

Just tell us the name of the nominee, the phone number or e-mail address at which the nominee can be reached, and, most important, why you feel the nominee deserves the award. You can attach supporting documentation. Please include your own name, address, and phone number.

We're looking for the Pioneers of the Electronic Frontier that have made and are making a difference. Thanks for helping us find them,

—The Electronic Frontier Foundation

Call for Papers

The USENIX Association and the University of Michigan's Center for Information Technology Integration are sponsoring a *Workshop on File Systems*, to be held on May 21st and 22nd, 1992, on the campus of the University of Michigan, in Ann Arbor.

Topics

The goal of the workshop is to bring together researchers and practitioners on all aspects of file systems, including but not limited to:

- File system performance measurement and models;
- WORM and other optical systems;
- Log-structured, RAID, and other high-performance systems;
- Mass-storage and archival systems;
- Support for replication, consistency, and mobility in distributed file systems;
- Naming and location in very-large distributed file systems.

Format

Workshop participants will be invited to present formal papers or informal "works-in-progress." There will be opportunities to identify and discuss trends, themes, and theoretical and practical aspects of file systems research and development. There will also be opportunities for "less polished" papers to be presented in the workshop sessions. Working sessions will consist of 15–20 submitted papers, to be presented in 20 and 30 minute time periods.

Paper submissions

You are invited to submit original papers from any area related to file systems for presentation at the workshop and inclusion in the proceedings. Several categories of papers will be considered:

- Original and unpublished research reports;
- Reports of innovative applications of current technology to new problem domains;
- Position papers on controversial points of practical or theoretical interest.

Five copies of a full paper or extended abstract should be sent to:

Workshop on File Systems
Center for Information Technology Integration
The University of Michigan
519 W. William Street
Ann Arbor, MI 48103-4943

Papers should include an attached separate front sheet describing the title of the paper, the name(s) of the author(s), affiliation, mailing address, telephone, fax, and e-mail address. Papers will be selected by the program committee based on originality, relevance, and impact.

Important dates

Manuscripts due: March 15, 1992
Program decision: April 1, 1992
Camera-ready copy due: April 15, 1992

For further information, contact:

fsworkshop@citi.umich.edu
+1 313-763-4403
+1 313-763-4434 (Fax)

Workshop Announcement

The *Joint EurOpen / USENIX Workshop on Portability and Interoperability* will be held April 6–9, 1992 in Jersey (Channel Islands, UK).

Goals One of the major goals of *Open Systems* is to support the distribution and mobility of data, users and processing. The purpose of this joint workshop is to identify the challenges of this goal and to discuss possible solutions. The workshop will address issues of concern to the system designer, developer, manager, administrator and end-user.

Topics Suggested topics for this workshop include:

System design and development issues:

- Software design, development and management
- Portability tips and techniques
- Standards? Help or hindrance?

System administration and management issues:

- How to cope with the old world
- Migration tools
- Resource sharing
- Public procurement
- Migration strategies

End user issues:

- Availability of tools and applications
- Education and training
- Why open systems? Why proprietary systems?
- Migrating from a proprietary to an open system environment

More information For more information contact:

EurOpen Secretariat
Owles Hall
Buntingford
Hertfordshire SG9 9PL
United Kingdom
Fax: +44 763 73255
europen@eu.net or europen-jersey@eu.net

Write to *ConneXions*!

Have a question about your subscription? Are you moving, and need to give us your new address? Suggestions for topics? Want to write an article? A letter to the Editor? Have a question for an author? Want to enquire about back issues? (there are now nearly sixty to choose from; ask for our free 1987–1991 index sheets). We want to hear from you. Simply write, call, fax, or e-mail us at:

ConneXions—The Interoperability Report
480 San Antonio Road, Suite 100
Mountain View, CA 94040–1219
USA
Phone: +1 415-941-3399 or 1-800-INTEROP (Toll-free in the USA)
Fax: +1 415-949-1779
E-mail: connexions@interop.com

CONNEXIONS

480 San Antonio Road
Suite 100
Mountain View, CA 94040
415-941-3399
FAX: 415-949-1779

FIRST CLASS MAIL
U.S. POSTAGE
PAID
SAN JOSE, CA
PERMIT NO. 1

ADDRESS CORRECTION
REQUESTED

CONNEXIONS

EDITOR and PUBLISHER Ole J. Jacobsen

EDITORIAL ADVISORY BOARD Dr. Vinton G. Cerf, Vice President,
Corporation for National Research Initiatives

A. Lyman Chapin, Chief Network Architect,
BBN Communications

Dr. David D. Clark, Senior Research Scientist,
Massachusetts Institute of Technology

Dr. David L. Mills, Professor,
University of Delaware

Dr. Jonathan B. Postel, Communications Division Director,
University of Southern California, Information Sciences Institute

Subscribe to CONNEXIONS

U.S./Canada ☐ \$150. for 12 issues/year ☐ \$270. for 24 issues/two years ☐ \$360. for 36 issues/three years

International \$ 50. additional per year (Please apply to all of the above.)

Name _____ Title _____

Company _____

Address _____

City _____ State _____ Zip _____

Country _____ Telephone () _____

☐ Check enclosed (in U.S. dollars made payable to CONNEXIONS).

☐ Visa ☐ MasterCard ☐ American Express ☐ Diners Club Card # _____ Exp. Date _____

Signature _____

Please return this application with payment to:

CONNEXIONS

Back issues available upon request \$15./each
Volume discounts available upon request

480 San Antonio Road, Suite 100
Mountain View, CA 94040 U.S.A.
415-941-3399 FAX: 415-949-1779

CONNEXIONS